Toward Better Gossip Simulation in Emergent Narrative Systems

Max Kreminski
Santa Clara University
mkreminski@scu.edu

***Abstract*—Interactive emergent narrative games often make use of social simulation techniques, including the modeling of character relationships and knowledge, to generate compelling gameplay and stories. Existing attempts at simulating the spread of knowledge and opinions between characters, however, strictly limit the informational and social content of the communications that characters exchange, making it difficult to create play experiences in which gossip is a core mechanic. In this paper, we introduce gossip simulation as an open problem in intelligent narrative technologies and present an abstract approach to informationally and socially rich gossip simulation for emergent narrative-oriented social simulations, as well as a preliminary concrete implementation of this approach.**

***Index Terms*—procedural content generation, emergent narrative, social simulation, story sifting**

I. Introduction

A number of interactive *emergent narrative* [1]–[3] games and play experiences make use of *social simulation* techniques [4], [5] to produce compelling stories "from the bottom up", via interactions between a cast of autonomous or semi-autonomous agents. Both commercially successful AAA and indie games (such as *The Sims* and *Dwarf Fortress* [6]) and high-profile research games (such as *Prom Week* [7]) adhere to this approach, and studies of player *retellings* of emergent narrative gameplay [8] have suggested that simulated social dynamics play a large role in helping players to frame their gameplay experiences as stories [9].

Several phenomena that have been deliberately simulated by emergent narrative game researchers and developers—including knowledge transfer between characters [10], character evaluation of other characters based on the actions they are known to have performed [11], and characters making sense of the world through narrative [12]—are bound together by their connection to *gossip*. Gossip is a prototypical category rather than a binary one, so there is no single bright line between gossip and non-gossip—but in general, gossip is typically viewed as a social phenomenon in which people tell one another implicitly or explicitly evaluative stories about other people's private lives [13]. It simultaneously serves information-sharing, interpersonal bond-forming and moral norm-establishing functions, and is ideally narrative in form. In addition, gossip is often considered inherently enjoyable [13] and is arguably the oldest form of narrative [14].

979-8-3503-2277-4/23/$31.00 ©2023 IEEE

Despite the apparent relatedness of gossip to many of the things that make social simulation games fun, however, no existing emergent narrative system that we are aware of is capable of full and systematic simulation of narrative gossip. Consequently, games and generative narrative systems that aim to implement key features of gossip—whether as game mechanics for players to interact with, elements of background simulation used to create more believable simulated characters, or as a narrative generation pipeline for gossipy stories that can be presented to a human interactor—have not yet reached their full potential.

In this paper, we briefly introduce gossip simulation as an open problem for intelligent narrative technologies research and extract some desiderata for gossip simulation from prior incomplete attempts. We then introduce a novel abstract approach to gossip simulation for emergent narrative-oriented social simulations, which is capable of achieving all of these desiderata, and a preliminary concrete implementation of this approach. We hope that this work paves the way for more sophisticated modeling of gossip in future interactive, emergent, and generative narrative systems.

II. Prior Work

In the context of emergent narrative-oriented social simulation, several attempts have been made to systematically simulate phenomena related to gossip. We argue, however, that none of these attempts make it possible for characters to exchange *informationally rich* and *narratively structured* accounts of what they *believe* other characters have *done*—though several systems have come close.

A. Rules-Based Approaches

To date, the majority of narrative-oriented social simulations have been based on hardcoded rules that govern character interactions. One prototypical system within this category is the Talk of the Town social simulation engine [10], which features character knowledge acquisition and propagation as a core feature of its simulation architecture. Each character has a set of *perceptible attributes* (including for instance what they look like, where they live, and their most recent whereabouts), and characters can both originate (for instance, by witnessing or confabulating) and propagate (by telling other characters) knowledge about these attributes. Every item of character knowledge is subject to potential mutation and decay, and when two characters meet, they decide by mutual

familiarity which other characters to exchange knowledge about. Altogether, this approach results in characters that can "observe, tell, misremember, and lie" [15] in sufficiently interesting ways to enable the core detective-work gameplay of Bad News [16], albeit mostly about things like a character's hair color and current place of residence—facts that mostly lack the informational, social, and narrative richness which makes gossip so enjoyable in real life [13].

The Versu interactive narrative framework [11] features gossip mechanisms that are in some ways very similar to our approach. Characters can evaluate one another with respect to various social *roles* (e.g., one character might judge another as a dutiful husband, an irresponsible brother, and a bad participant at the ball); communicate these evaluations, and the reasons for these evaluations, to other characters; and choose to adopt another character's communicated evaluation. However, Versu permits the storage of only one evaluation and justification per character-character-role triple, strictly limiting the complexity of the evaluations that characters can hold and communicate. This may be due to Versu's focus on fine-grained, beat-by-beat simulation of dialogue among a relatively small cast of characters: in such a gameplay context, there is little room for characters to gradually assemble and eventually communicate multi-part stories about one another's lives, so full-fledged simulation of narrative gossip is not needed to support the system's goals. Nevertheless, Versu's gossip simulation is still capable of producing interesting moments of emergent narrative—for instance, when one of the system's developers found (to his surprise) that he had unintentionally alienated several characters he had not yet encountered by behaving rudely toward a butler, who then gossiped about the player character's rudeness behind the scenes [17].

In the Talk of the Town-inspired Kismet social simulation engine [18], characters can gossip about individual character actions that they have witnessed or been told about by other characters. Knowledge of past actions can then influence how characters behave toward one another, including via inference rules that trigger only when a single character comes to know about multiple different past actions by a single other character (and, implicitly, enframing these actions as a more complicated story about that character). However, Kismet does not seem to support the *communication* between characters of these more complicated stories, resulting in significant loss of informational and evaluative structure when characters communicate about one another's actions.

Finally, in the social and management simulation game *City of Gangsters* [19], each character (including the player character) accumulates a separate reputation with each other character. This reputation is based on specific actions that a character has performed (mostly various forms of helping or hurting others); each action is automatically known to its direct target, but can also be communicated by its target to other NPCs that are directly connected to the target in the social network. Consequently, characters can be said to engage in gossip with their direct social connections about specific actions that other characters have performed—though here again, gossip communicates only one action at a time; the evaluation attached to each action is very simple in nature (either a positive or negative integer score); and it remains unclear whether a single piece of gossip can be passed on further by a character who has merely *heard about* (rather than witnessed directly) the event that it involves.

B. Language Model-Based Approaches

One promising and very recently introduced alternative approach to social simulation [20] makes use of a large language model (LLM) to simulate characters that plan, remember, and engage in dialogue with one another. Early experiments with this approach seem to enable characters to share open-ended natural language summaries of what they know about other characters, which the recipient of the summary can then interpret to form new memories, and these new memories can then be passed along to other characters in the same way. The shared stories sometimes have both informational and evaluative components, though they don't seem to exhibit high levels of complexity and they don't seem to take narrative form: example interactions presented by Park et al. [20] generally communicate only a single action or judgment at a time. Though the exact limitations of this approach are difficult to evaluate due to ambiguities in the apparent capabilities of the underlying LLM, this approach nevertheless seems to present a high-level simulation architecture that could eventually prove capable of implementing the full range of gossip dynamics discussed in the sociological literature.

C. Summary of Prior Work

As far as we are aware, no existing approach to gossip simulation simultaneously (A) allows characters to gossip about *character actions* rather than basic facts; (B) enframes these actions as *microstories* that can communicate arbitrarily complicated *justified perceptions* of characters; and (C) allows other characters to granularly *reinterpret* the actions they've learned about, so that the same piece of gossip can cause different reactions in different characters based on what else each recipient believes about the gossip's subject. (The LLM-driven approach presented by Park et al. [20] may be an exception, but the full capabilities of this approach are difficult to gauge.) Consequently, it is this combination of features—which seem necessary to implement the gossip dynamics observed in sociological studies—that we aim to include in our new model of gossip.

III. Gossip Generation and Propagation

In order to unify the gossip-related features of prior architectures under a single model that is capable of implementing the full range of gossip dynamics discussed in the sociological literature, a novel abstract approach to gossip simulation seems to be needed. At a high level, our proposed approach to gossip generation and propagation has four main phases:

- The **witness phase**. Whenever a character takes an action, we must determine which (if any) other characters witnessed the action; how (if at all) to *mutate* their perception

of the action, allowing for possible misunderstanding of witnessed actions; whether (if at all) each character should *record* their perception of the action; and how strong each character's recorded perception of the action (if any) should initially be.
- The **reflection phase**. Every character that participates in gossip simulation must occasionally *reflect on* the actions they remember, assembling these memories into *microstories* that can be communicated to other characters.
- The **propagation phase**. When two characters interact, we must determine which (if any) microstories each character will share with the other; how (if at all) to *mutate* the action memories underlying shared stories in their transmission, allowing for possible misunderstanding of communicated actions; and how strong the receiving character's recorded perception of any communicated stories or actions should initially be.
- The **decay phase**. Every character that participates in gossip simulation must occasionally *decay* their memories of actions and stories, decreasing the *weight* value of each remembered entity by a certain amount and culling any memories for which the weight value falls below a fixed threshold. This ensures that recent actions are prioritized as gossip over older actions by default, and that characters do not accumulate larger and larger numbers of memories as the simulation continues to run.

To participate in gossip simulation, a character must have a **memory database**. This database stores a character's memories of both individual actions and higher-level microstories involving these actions.

Optionally, characters' other behaviors can also be adjusted based on their impressions of one another. How exactly to do this is outside the scope of this work, but one simple initial approach might involve associating every microstory with a mapping of involved character roles to numerical liking scores. The recipient of a story would then iterate over the character roles involved in the story and add the appropriate amount to their overall liking of the character playing that role. Finally, these cumulative liking scores could be used to influence characters' behaviors toward one another in appropriate ways.

IV. Concrete Implementation Details

In addition to our proposed abstract architecture for gossip simulation, we also introduce Gossamer: a preliminary and open-source[1] concrete implementation of this architecture, intended to illustrate how our abstract approach can be realized in (JavaScript) code. Gossamer is essentially a computational caricature [21] of gossip behaviors, intended to embody our claims that a full-featured gossip simulation requires all four of the phases outlined in the previous section.

Beyond straightforward implementations of the previously described phases, Gossamer also contains two key pieces of key infrastructure that would likely be useful in future attempts at implementing similar systems: a per-character memory database system and a procedure for determining the salience of storyworld entities to specific characters.

A. Character Memory Databases

In our initial JavaScript implementation of Gossamer, every character stores their memories of both actions and stories in a character-specific DataScript database[2]. This allows for easy querying of memories via the Felt story sifting language [22] during the reflection phase and improves query performance by ensuring that these queries are implicitly able to ignore any information that the reflecting character doesn't know.

The use of Felt-style story sifting for memory queries, meanwhile, may enable future implementation of more sophisticated story sifting heuristics—such as the StU heuristic, which prioritizes statistically unlikely stories over likelier ones and correlates well with human judgments of story interestingness [23]—as a means of selecting which microstories characters will remember and share.

B. Salience Determination

At every phase of gossip simulation, it seems important for a simulator to be able to determine a character's level of familiarity with or interest in other storyworld entities, including characters and events. During the witness phase, a model of familiarity or interest can be used to determine both which witnessed events a character records and the likelihood that these events will be mutated during storage. During the reflection phase, this model can be used to determine which events to enframe as microstories that influence the behavior of the reflecting character and can potentially be propagated later on. During the propagation phase, the microstories to be shared between a given pair of characters might be influenced by both characters' familiarity with or interest in the topics of those microstories. And during the decay phase, familiarity or interest can be used to determine which memories decay more or less rapidly.

To meet these needs, we introduce a model of entity *salience* as part of our gossip implementation. Salience models have been applied in prior intelligent narrative technologies research to improve the humanlikeness of computational systems, for instance in the Indexter model of storyfulness [24]. By modeling the extent to which specific events, characters, and other storyworld entities are salient to a specific character, we enable probabilistic reasoning about how likely a character is to form memories of witnessed events; to remember witnessed events *correctly*, rather than in some mutated form; to communicate stories about these events to other characters; and to preserve, rather than rapidly forgetting, their memories of these events.

Specifically, we model salience in terms of *familiarity*: the salience of a storyworld entity to a particular character is equal to the number of times that entity is referred to by other entities in the character's memory database, normalized by the total number of entity references in the memory database as a whole. Though this is a very rough proxy for salience, it has

[1] https://github.com/mkremins/gossamer

[2] https://github.com/tonsky/datascript

several desirable effects: in particular, it increases the tendency of characters to gossip about characters they regularly interact with or hear about from other characters, as well as for characters to formulate and share *self-centered* stories (since characters are highly familiar, and thus highly salient, to themselves almost by definition). Improvements to this simple salience model may be introduced in future work.

V. Conclusion and Future Work

We have introduced gossip simulation as an open problem in intelligent narrative technologies research; surveyed prior work to extract a set of preliminary desiderata for gossip simulation; and introduced an abstract architecture for a gossip model that seems capable of meeting these desiderata, as well as a preliminary concrete implementation of this model (Gossamer). To date, we have not yet conducted much evaluation of our implementation, and evaluation is thus a top priority for future work. In particular, we plan to evaluate our model by testing whether certain emergent gossip-related phenomena appear. These phenomena include, but are not limited to:

- Rumors spreading rapidly within well-connected subgraphs of the social network, but taking longer to diffuse across the full graph [25]
- Rumors being displaced by rumors about more shocking events, when the latter are deliberately introduced
- Characters being preceded by their reputation [25]
- Perception of certain characters becoming polarized due to "dueling rumors" about their actions
- Certain characters (who are deliberately made widely known to other characters) becoming celebrities who are "famous for being famous", due to their frequent selection as a topic of gossip [26]
- Characters with power over others becoming more frequent topics of gossip, especially for those characters they have power over (e.g., a boss to their employees) [26]

In addition, we also hope to integrate our approach into a research-scale interactive emergent narrative game and gauge player experience via standard playtesting methods. This will be a necessary step for understanding the range of play experiences that our gossip model supports and its feasibility for integration into larger games.

References

[1] S. Louchart, J. Truesdale, N. Suttie, and R. Aylett, "Emergent narrative: past, present and future of an interactive storytelling approach," in *Interactive Digital Narrative*, pp. 185–199, Routledge, 2015.

[2] R. Walsh, "Emergent narrative in interactive media," *Narrative*, vol. 19, no. 1, pp. 72–85, 2011.

[3] M. Kreminski and M. Mateas, "A coauthorship-centric history of interactive emergent narrative," in *Interactive Storytelling: 14th International Conference on Interactive Digital Storytelling, ICIDS 2021*, pp. 222–235, Springer, 2021.

[4] S. Azad and C. Martens, "Little computer people: A survey and taxonomy of simulated models of social interaction," in *Proceedings of the ACM on Human-Computer Interaction*, vol. 5, ACM New York, NY, USA, 2021.

[5] S. Johnson-Bey, M. J. Nelson, and M. Mateas, "Exploring the design space of social physics engines in games," in *Interactive Storytelling: 15th International Conference on Interactive Digital Storytelling, ICIDS 2022*, pp. 559–576, Springer, 2022.

[6] T. Adams, "Emergent narrative in Dwarf Fortress," in *Procedural Storytelling in Game Design*, pp. 149–158, AK Peters/CRC Press, 2019.

[7] J. McCoy, M. Treanor, B. Samuel, A. A. Reed, M. Mateas, and N. Wardrip-Fruin, "Social story worlds with Comme il Faut," *IEEE Transactions on Computational Intelligence and AI in Games*, vol. 6, no. 2, pp. 97–112, 2014.

[8] M. P. Eladhari, "Re-tellings: the fourth layer of narrative as an instrument for critique," in *Interactive Storytelling: 11th International Conference on Interactive Digital Storytelling, ICIDS 2018*, pp. 65–78, Springer, 2018.

[9] M. Kreminski, B. Samuel, E. Melcer, and N. Wardrip-Fruin, "Evaluating AI-based games through retellings," in *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, vol. 15, pp. 45–51, 2019.

[10] J. Ryan and M. Mateas, "Simulating character knowledge phenomena in Talk of the Town," in *Game AI Pro 360*, pp. 135–150, CRC Press, 2019.

[11] R. Evans and E. Short, "Versu—a simulationist storytelling system," *IEEE Transactions on Computational Intelligence and AI in Games*, vol. 6, no. 2, pp. 113–130, 2013.

[12] M. Kreminski, M. Dickinson, M. Mateas, and N. Wardrip-Fruin, "Why Are We Like This?: Exploring writing mechanics for an AI-augmented storytelling game," in *Proceedings of the Electronic Literature Organization Conference*, 2020.

[13] A. Ben-Ze'ev, "The vindication of gossip," in *Good Gossip* (R. F. Goodman and A. Ben-Ze'ev, eds.), University Press of Kansas, 1994.

[14] M.-L. Ryan, "From narrative games to playable stories: Toward a poetics of interactive narrative," *StoryWorlds: a journal of narrative studies*, vol. 1, pp. 43–59, 2009.

[15] J. Ryan, A. Summerville, M. Mateas, and N. Wardrip-Fruin, "Toward characters who observe, tell, misremember, and lie," in *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, vol. 11, pp. 56–62, 2015.

[16] B. Samuel, J. Ryan, A. J. Summerville, M. Mateas, and N. Wardrip-Fruin, "Bad News: An experiment in computationally assisted performance," in *Interactive Storytelling: 9th International Conference on Interactive Digital Storytelling, ICIDS 2016*, pp. 108–120, Springer, 2016.

[17] E. Short and R. Evans, "Versu: Conversation implementation." https://emshort.blog/2013/02/26/versu-conversation-implementation/, 2013.

[18] A. Summerville and B. Samuel, "Kismet: a small social simulation language," in *ICCC-WS 2020: Proceedings of the Joint Workshops of the International Conference on Computational Creativity*, 2020.

[19] R. Zubek, I. Horswill, E. Robison, and M. Viglione, "Social modeling via logic programming in City of Gangsters," in *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, vol. 17, pp. 220–226, 2021.

[20] J. S. Park, J. C. O'Brien, C. J. Cai, M. R. Morris, P. Liang, and M. S. Bernstein, "Generative agents: Interactive simulacra of human behavior," *arXiv preprint arXiv:2304.03442*, 2023.

[21] A. Smith and M. Mateas, "Computational caricatures: Probing the game design process with AI," in *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, vol. 7, pp. 14–18, 2011.

[22] M. Kreminski, M. Dickinson, and N. Wardrip-Fruin, "Felt: a simple story sifter," in *Interactive Storytelling: 12th International Conference on Interactive Digital Storytelling, ICIDS 2019*, pp. 267–281, Springer, 2019.

[23] M. Kreminski, M. Dickinson, N. Wardrip-Fruin, and M. Mateas, "Select the unexpected: A statistical heuristic for story sifting," in *Interactive Storytelling: 15th International Conference on Interactive Digital Storytelling, ICIDS 2022*, pp. 292–308, Springer, 2022.

[24] R. E. Cardona-Rivera, K. B. Cassell, S. G. Ware, and R. M. Young, "Indexter: A computational model of the event-indexing situation model for characterizing narratives," in *Proceedings of the 3rd Workshop on Computational Models of Narrative*, pp. 34–43, 2012.

[25] N. Emler, "Gossip, reputation, and social adaptation," in *Good Gossip* (R. F. Goodman and A. Ben-Ze'ev, eds.), University Press of Kansas, 1994.

[26] M. Ayim, "Knowledge through the grapevine: Gossip as inquiry," in *Good Gossip* (R. F. Goodman and A. Ben-Ze'ev, eds.), University Press of Kansas, 1994.